arXiv:1009.3214v2 [cs.SC] 1 Jan 2011

Computing sparse multiples of polynomials*

Mark Giesbrecht, Daniel S. Roche, and Hrushikesh Tilak
Cheriton School of Computer Science, University of Waterloo

January 4, 2011

Abstract

We consider the problem of finding a sparse multiple of a polynomial. Given $f \in \mathsf{F}[x]$ of degree d over a field F, and a desired sparsity t, our goal is to determine if there exists a multiple $h \in \mathsf{F}[x]$ of f such that h has at most t non-zero terms, and if so, to find such an h. When $\mathsf{F} = \mathbb{Q}$ and t is constant, we give a polynomial-time algorithm in d and the size of coefficients in h. When F is a finite field, we show that the problem is at least as hard as determining the multiplicative order of elements in an extension field of F (a problem thought to have complexity similar to that of factoring integers), and this lower bound is tight when $t = 2$.

1 Introduction

Let F be a field, which will later be specified either to be the rational numbers ($\mathbb{Q}$) or a finite field with q elements ($\mathbb{F}_q$). We say a polynomial $h \in \mathsf{F}[x]$ is *t-sparse* (or *has sparsity t*) if it has at most t nonzero coefficients in the standard power basis; that is, h can be written in the form

$$h = h_1 x^{d_1} + h_2 x^{d_2} + \ldots + h_t x^{d_t} \quad \text{for } h_1, \ldots, h_t \in \mathsf{F} \text{ and } d_1, \ldots, d_t \in \mathbb{N}. \quad (1.1)$$

Sparse polynomials have a compact representation as a sequence of coefficient-degree pairs $(h_1, d_1), \ldots, (h_t, d_t)$, which allow representation and manipulation of very high degree polynomials. Let $f \in \mathsf{F}[x]$ have degree d. We examine the computation of a t-sparse multiple of f. That is, we wish to determine if there exist $g, h \in \mathsf{F}[x]$ such that $fg = h$ and h has prescribed sparsity t, and if so, to find such an h. We do not attempt to find g, as it may have a super-polynomial number of terms, even though h has a compact representation (see Theorem 3.7).

Sparse multiples over finite fields have cryptographic applications. Their computation is used in correlation attacks on LFSR-based stream ciphers (El Aimani and von zur Gathen, 2007; Didier and Laigle-Chapuy, 2007). The

*The authors would like to thank the Natural Sciences and Engineering Research Council of Canada (NSERC), and MITACS

security of the TCHo cryptosystem is also based on the conjectured computational hardness of sparsest multiple computation over $\mathbb{F}_2[x]$ (Aumasson et al., 2007); our results provide further evidence that this is in fact a computationally difficult problem.

Sparse multiples can facilitate efficient arithmetic in extension fields (Brent and Zimmermann, 2003) and in designing interleavers for error-correcting codes (Sadjadpour et al., 2001). The linear algebra formulation in Section 2 relates to finding the minimum distance of a binary linear code (Berlekamp et al., 1978; Vardy, 1997) as well as finding "sparsifications" of linear systems (Egner and Minkwitz, 1998).

One of our original motivations was to understand the complexity of sparse polynomial *implicitization* over $\mathbb{Q}$ or $\mathbb{R}$: Given a curve represented explicitly as a set of parametric rational functions, find a sparse polynomial whose zero set contains all points on the curve (see, e.g., Emiris and Kotsireas (2005)). This is a useful operation in computer aided geometric design for facilitating various operations on the curve, and work here can be thought of as a univariate version of this problem.

We often consider the related problem of finding a sparse annihilator for a set of points — that is, a sparse polynomial with given roots. This is exactly equivalent to our problem when the input polynomial f is squarefree, and in the binomial case corresponds to asking whether a given root can be written as a surd. This is also the problem we are really interested in regarding implicitization, and allows us to build on significant literature from the number theory community on the roots of sparse polynomials.

In general, we assume that the desired sparsity t is a constant. This seems reasonable given that over a finite field, even for $t = 2$, the problem is probably computationally hard (Theorem 5.1). In fact, we have reason to conjecture that the problem is intractable over $\mathbb{Q}$ or $\mathbb{F}_q$ when t is a parameter. Our algorithms are exponential in t but polynomial in the other input parameters when t is constant.

Over $\mathbb{Q}[x]$, the analysis must consider coefficient size, and we will count machine word operations in our algorithms to account for coefficient growth. We follow the conventions of Lenstra (1999) and define the *height* of a polynomial as follows. Let $f \in \mathbb{Q}[x]$ and $r \in \mathbb{Q}$ the least positive rational number such that $rf \in \mathbb{Z}[x]$. If $rf = \sum_i a_i x^{e_i}$ with each $a_i \in \mathbb{Z}$, then the *height* of f, written $\mathcal{H}(f)$, is $\max_i |a_i|$.

We examine variants of the sparse multiple problem over $\mathbb{F}_q$ and $\mathbb{Q}$. Since every polynomial in $\mathbb{F}_q$ has a 2-sparse multiple of high degree, given $f \in \mathbb{F}_q[x]$ and $n \in \mathbb{N}$ we consider the problem of finding a t-sparse multiple of f with degree at most n. For input $f \in \mathbb{Q}[x]$ of degree d, we consider algorithms which seek t-sparse multiples of height bounded above by an additional input value $c \in \mathbb{N}$. We present algorithms requiring time polynomial in d and $\log c$.

The remainder of the paper is structured as follows.

In Section 2, we consider the straightforward linear algebra formulation of the sparse multiple problem. This is useful over $\mathbb{Q}[x]$ once a bound on the output degree is derived, and also allows us to bound the output size. In addition, it

connects our problems with related NP-complete coding theory problems.

In Section 3 we consider the problem of finding the least-degree binomial multiple of a rational polynomial. A polynomial-time algorithm in the size of the input is given which completely resolves the question in this case. This works despite the fact that we show polynomials with binomial multiples whose degrees and heights are both exponential in the input size!

In Section 4 we consider the more general problem of finding a t-sparse multiple of an input $f \in \mathbb{Q}[x]$. Given a height bound $c \in \mathbb{N}$ we present an algorithm which requires polynomial time in $\deg f$ and $\log c$, except in the very special case that f has both non-cyclotomic and repeated cyclotomic factors.

Section 5 shows that, even for $t = 2$, finding a t-sparse multiple of a polynomial $f \in \mathbb{F}_q[x]$ is at least as hard as finding multiplicative orders in an extension of $\mathbb{F}_q$ (a problem thought to be computationally difficult). This lower bound is shown to be tight for $t = 2$ due to an algorithm for computing binomial multiples that uses order finding.

Open questions and avenues for future research are discussed in Section 6.

An extended abstract of some of this work appears in Giesbrecht, Roche, and Tilak (2010). Some of this work and further explorations, also appears in the Masters thesis of Tilak (2010).

2 Linear algebra formulation

The sparsest multiple problem can be formulated using linear algebra. This requires specifying bounds on degree, height and sparsity; later some of these parameters will be otherwise determined. This approach also highlights the connection to some problems from coding theory. We exhibit a randomized algorithm for finding a t-sparse multiple h of a degree-d polynomial $f \in \mathbb{Q}[x]$, given bounds c and n on the height and degree of the multiple respectively. When t is a constant, the algorithm runs in time polynomial in n and $\log \mathcal{H}(f)$ and returns the desired output with high probability. We also conjecture the intractability of some of these problems, based on similar problems in coding theory. Finally, we show that the construction of Vardy (1997) can be used to show the problem of finding the sparsest vector in an integer lattice is NP-complete, which was conjectured by Egner and Minkwitz (1998).

Let R be a principal ideal domain, with $f \in \mathsf{R}[x]$ of degree d and $n \in \mathbb{N}$ given. Suppose $g, h \in \mathsf{R}[x]$ have degrees $n - d$ and n respectively, with $f = \sum_0^d f_i x^i$, $g = \sum_0^{n-d} g_i x^i$ and $h = \sum_0^n h_i x^i$. The coefficients in the equation $fg = h$ satisfy

the following linear system:

$$\underbrace{\begin{bmatrix} f_0 & & & \\ f_1 & f_0 & & \\ \vdots & f_1 & \ddots & \\ f_d & \vdots & \ddots & f_0 \\ & f_d & \ddots & f_1 \\ & & \ddots & \vdots \\ & & & f_d \end{bmatrix}}_{A_{f,n}} \underbrace{\begin{bmatrix} g_0 \\ g_1 \\ \\ \vdots \\ \\ g_{n-d} \end{bmatrix}}_{v_g} = \underbrace{\begin{bmatrix} h_0 \\ h_1 \\ \\ \vdots \\ \\ h_n \end{bmatrix}}_{v_h}. \tag{2.1}$$

Thus, a multiple of f of degree at most n and sparsity at most t corresponds to a vector with at most t nonzero entries (i.e., a t-sparse vector) in the linear span of $A_{f,n}$.

If $f \in \mathsf{R}[x]$ is squarefree and has roots $\{\alpha_1, \ldots, \alpha_d\}$, possibly over a finite extension of R, then the following also holds:

$$\underbrace{\begin{bmatrix} 1 & \alpha_1 & \cdots & \alpha_1^n \\ 1 & \alpha_2 & \cdots & \alpha_2^n \\ \vdots & \vdots & \vdots & \vdots \\ 1 & \alpha_d & \cdots & \alpha_d^n \end{bmatrix}}_{A_n(\alpha_1,\ldots,\alpha_d)} \begin{bmatrix} h_0 \\ h_1 \\ \\ \vdots \\ \\ h_n \end{bmatrix} = \mathbf{0}. \tag{2.2}$$

Thus t-sparse multiples of a squarefree f correspond to t-sparse R-vectors in the nullspace of $A_n(\alpha_1, \ldots, \alpha_d)$.

2.1 Finding short l_∞ vectors in lattices

This technical section presents a randomized, polynomial-time algorithm to find the shortest l_∞ vector in a constant-dimensional lattice. Our algorithm is a modification of Ajtai et al. (2001), based on the presentation by Regev (2004), adapted to the case of infinity norm. Since this the techniques are essentially drawn from the literature, and while necessary, are not the central thrust of this current paper, full details are left to Appendix A.

Algorithm 2.1 below starts by computing a rough approximation of the shortest l_2 vector using LLL (Lenstra, Lenstra, and Lovász, 1982), and then scales the lattice accordingly. The main while loop then consists of two phases: sampling and sieving. First, a large number of random vectors $\{x_1, \ldots, x_m\}$ are sampled in an appropriately-sized ball around the origin. We take these modulo the basis B to obtain vectors $\{y_1, \ldots, y_m\}$ with the property that each $x_i - y_i$ is in the lattice of B. Next, we use a series of sieving steps in the while loop in Step 13 to find a small subset of the y_i vectors that are close to every other vector and

use these as "pivots". The pivots are discarded from the set, but all remaining lattice vectors $x_i - y_i$ are made smaller. After this, the set W_γ contains most lattice vectors whose l_2 length is close to γ.

Algorithm 2.1: Shortest l_∞ vector in a lattice

Input: Basis $U \in \mathbb{Z}^{n \times d}$ for an integer lattice $\mathcal{L}$ of dimension n and size $d \leq n$

Output: Shortest l_∞ vector in $\mathcal{L}$

1 $\lambda \leftarrow$ approximate l_2-shortest vector in $\mathcal{L}$ from Lenstra et al. (1982)
2 $B \leftarrow (1/\|\lambda\|_2) \cdot U$, stored as a list of vectors $[b_1, \ldots, b_d]$
3 **for** $k \in \{1, 2, \ldots, 2n\}$ **do**
4 $B \leftarrow 1.5 \cdot B$
5 $r_0 \leftarrow n \max_i \|b_i\|_2$
6 $\gamma \leftarrow 3/2$
7 **while** $\gamma \leq 3\sqrt{n} + 1$ **do**
8 $m \leftarrow \left\lceil 2^{(7+\lceil \log \gamma \rceil)n} \log r_0 \right\rceil$
9 Sample points $\{x_1, \ldots, x_m\}$ uniformly and independently from $\mathbf{B}_n(0, \gamma)$, the n-dimensional ball of radius γ centered around $\mathbf{0}$
10 $S \leftarrow \{1, 2, \ldots, m\}$
11 $y_i \leftarrow x_i \mod \mathcal{P}(B)$ for every $i \in S$, $\mathcal{P}(B)$ being the parallelogram of B defined in the proof of Lemma A.2
12 $r \leftarrow r_0$
13 **while** $r > 2\gamma + 1$ **do**
14 $J \leftarrow \emptyset$
15 **for** $i \in S$ **do**
16 **if** $\exists j \in J$ *such that* $\|y_j - y_i\| \leq r/2$ **then** $\eta_i \leftarrow j$
17 **else** $J \leftarrow J \cup \{i\}$
18 $S \leftarrow S \setminus J$
19 $y_i \leftarrow y_i + x_{\eta_i} - y_{\eta_i}$ for $i \in S$
20 $r \leftarrow r/2 + \gamma$
21 $Y_\gamma \leftarrow \{(x_i - y_i) \mid i \in S\}$
22 $W_\gamma \leftarrow \{v - w \mid v, w \in Y_\gamma \text{ and } v \neq w\}$
23 $\gamma \leftarrow 3\gamma/2$
24 $v_k \leftarrow$ shortest l_∞ vector in any W_γ
25 **return** *shortest* l_∞ *vector in* $\{(\|\lambda\|_2 / 1.5^k) \cdot v_k \mid k = 1, 2, \ldots, n\}$

If we are fortunate enough that the shortest l_2 vector in the lattice with basis B set on Step 4 has length between 2 and 3, then we know that the shortest l_∞ vector in this lattice must have l_2 length between 2 and $3\sqrt{n}$. By iterating γ in the appropriate range, we will encounter this shortest l_∞ vector and set it to v_k on Step 24 with high probability. We prove, given our approximate starting point from LLL, we will be in this "fortunate" situation in at least one iteration through the outer for loop.

The correctness and efficiency of the algorithm is given by the following theorem, whose proof we defer to Appendix A.

Theorem 2.1. *Given a lattice basis $U \in \mathbb{Z}^{n\times d}$, Algorithm 2.1 returns the shortest l_∞ vector in the lattice of U, with probability at least $1-1/2^{O(n)}$, using $2^{O(n\log n)}\cdot\|U\|^{O(1)}$ bit operations.*

2.2 Finding a sparse multiple of bounded height and degree

We now present an algorithm to find the sparsest bounded-degree, bounded-height multiple $h \in \mathbb{Q}[x]$ of an input $f \in \mathbb{Q}[x]$. Since $\mathcal{H}$ is invariant under scaling, we may assume that $f, g, h \in \mathbb{Z}[x]$.

The basic idea is the following. Having fixed the positions at which the multiple h has nonzero coefficients, finding a low-height multiple is reduced to finding the nonzero vector with smallest l_∞ norm in the image of a small lattice.

Let $I = \{i_1, \ldots, i_t\}$ be a t-subset of $\{0, \ldots, n\}$, and $A^I_{f,n} \in \mathbb{Z}^{(n-t+1)\times(n-d+1)}$ the matrix $A_{f,n}$ with rows $i_1, \ldots, i_t$ removed. Denote by $B^I_{f,n} \in \mathbb{Z}^{t\times(n-d+1)}$ the matrix consisting of the removed rows $i_1, \ldots, i_t$ of the matrix $A_{f,n}$. Existence of a t-sparse multiple $h = h_{i_1}x^{i_1} + h_{i_2}x^{i_2} + \cdots + h_{i_t}x^{i_t}$ of input f is equivalent to the existence of a vector v_g such that $A^I_{f,n}\cdot v_g = \mathbf{0}$ and $B^I_{f,n}\cdot v_g = [h_{i_1}, \ldots, h_{i_t}]^T$.

Now let $C^I_{f,n}$ be a matrix whose columns span the nullspace of the matrix $A^I_{f,n}$. Since $A_{f,n}$ has full column rank, the nullspace of $A^I_{f,n}$ has dimension $s \le t$ and $C^I_{f,n} \in \mathbb{Z}^{(n-d+1)\times s}$. Thus, a t-sparse multiple $h = h_{i_1}x^{i_1} + \cdots + h_{i_t}x^{i_t}$ of f exists if and only if there exists a $v \in \mathbb{Z}^s$ such that

$$B^I_{f,n}\cdot C^I_{f,n}\cdot v = [h_{i_1}, \ldots, h_{i_t}]^T. \tag{2.3}$$

Note that $B^I_{f,n}\cdot C^I_{f,n} \in \mathbb{Z}^{t\times s}$. Our approach, outlined in Algorithm 2.2, is to generate this lattice and search for a small, t-sparse vector in it. For completeness, we first define the subset ordering used in the search.

Definition 2.2. Let $a = (a_1, a_2, \ldots, a_k)$ and $b = (b_1, b_2, \ldots, b_k)$ be two k-tuples. a precedes b in *reverse lexicographical order* if and only if there exists an index i with $1 \le i \le k$ such that $a_i < b_i$, and for all j with $i < j \le k$, $a_j = b_j$.

The following lemma shows how to compute Step 5 efficiently using the Smith normal form.

Lemma 2.3. *Given $T \in \mathbb{Z}^{k\times\ell}$ with $k \ge \ell$ and nullspace of dimension s, we can compute a $V \in \mathbb{Z}^{\ell\times s}$ such that the image of V equals the nullspace of T. The algorithm requires $\tilde{O}(k\ell^2 s\log\|T\|)$ bit operations (ignoring logarithmic factors).*

Proof. First compute the Smith normal form of the matrix: $T = PSQ$ for diagonal matrix $S = \mathrm{diag}(\delta_1, \ldots, \delta_{\ell-s}, 0, \ldots, 0) \in \mathbb{Z}^{k\times\ell}$ and unimodular matrices $P \in \mathbb{Z}^{k\times k}$ and $Q \in \mathbb{Z}^{\ell\times\ell}$. Storjohann (2000) gives efficient algorithms to compute such a P, S, Q with $\tilde{O}(k\ell^2 s\log\|T\|)$ bit operations.

Then since any vector $\mathbf{v}$ in the nullspace of T satisfies $PSQ\mathbf{v} = \mathbf{0}$, $SQ\mathbf{v} = \mathbf{0}$ also and $\mathbf{v}$ is in the nullspace of SQ. Next compute the inverse of Q; this can

Algorithm 2.2: Bounded-Degree Bounded-Height Sparsest Multiple

Input: $f \in \mathbb{Z}[x]$ and $t, n, c \in \mathbb{N}$
Output: A t-sparse multiple $h \in \mathbb{Z}[x]$ of f with $\deg(h) \leq n$ and $\mathcal{H}(h) \leq c$, or "NONE"

1 **for** $s = 2, 3, \ldots, t$ **do**
2 **foreach** *s-subset $I = (0, i_2, \ldots, i_s)$ of $\{0, 1, \ldots, n\}$, sorted in reverse lexicographic order,* **do**
3 Compute matrices $A^I_{f,n}$ and $B^I_{f,n}$ as defined above
4 **if** $A^I_{f,n}$ *does not have full column rank* **then**
5 Compute matrix $C^I_{f,n}$, a kernel basis for $A^I_{f,n}$
6 $\mathbf{h} \leftarrow$ shortest l_∞ vector in the lattice of $B^I_{f,n} \cdot C^I_{f,n}$ from Algorithm 2.1
7 **if** $\|\mathbf{h}\|_\infty \leq c$ **then return** $h_1 + h_2 x^{i_2} + \cdots + h_t x^{i_t}$
8 **return** "NONE"

be accomplished with the same number of bit operations since $\ell \leq k$. Define V to be the last s columns of Q^{-1}. Due to the diagonal structure of S, V must be a nullspace basis for SQ, and furthermore V has integer entries since Q is unimodular. □

The correctness and efficiency of Algorithm 2.2 can then be summarized as follows.

Theorem 2.4. *Algorithm 2.2 correctly computes a t-sparse multiple h of f of degree n and height c, if it exists, with $(\log \mathcal{H}(f))^{O(1)} \cdot n^{O(t)} \cdot 2^{O(t \log t)}$ bit operations. The sparsity s of h is minimal over all multiples with degree less than n and height less than c, and the degree of h is minimal over all such s-sparse multiples.*

Proof. The total number of iterations of the for loops is $\sum_{s=2}^{t} \binom{n-1}{s-1} < n^t$. Computing the rank of $A^I_{f,n}$, and computing the matrices $B^I_{f,n}$ and $C^I_{f,n}$ can each be done in polynomial time by Lemma 2.3. The size of the entries of $C^I_{f,n}$ is bounded by some polynomial $(\log \mathcal{H}(h) + n)^{O(1)}$. The computation of the shortest l_∞ vector can be done using $2^{O(t \log t)}$ operations on numbers of length $(\log \mathcal{H}(h) + n)^{O(1)}$, by Theorem 2.1.

The minimality of sparsity and degree comes from the ordering of the for loops. Specifically, the selection of subsets in Step 2 is performed in *reverse lexicographic order*, so that column subsets I corresponding to lower degrees are always searched first. □

2.3 Relationship to NP-hard problems

Note that the above algorithms require time exponential in t, and are only polynomial-time for constant t. It is natural to ask whether there are efficient

algorithms which require time polynomial in t. We conjecture this problem is probably NP-complete, and point out two results of Vardy (1997) and Guruswami and Vardy (2005) on related problems that are known to be hard.

The formulation (2.2) seeks the sparsest vector in the nullspace of a (structured) matrix. For an unstructured matrix over finite fields, this is the problem of finding the minimum distance of a linear code, shown by Vardy (1997) to be NP-complete. The same problem over integers translates into finding the sparsest vector in an integer lattice. It was posed as an open problem in Egner and Minkwitz (1998). Techniques similar to Vardy (1997) prove that this problem is also NP-complete over the integers, a fact proved in Theorem 2.5.

Of course, the problem may be easier for structured matrices as in (2.2) However, Guruswami and Vardy (2005) show that maximum likelihood decoding of cyclic codes, which seeks sparse solutions to systems of equations of similar structure to (2.2), is also NP complete. They do require the freedom to choose a right-hand-side vector, whereas we insist on a sparse vector in the nullspace. While these two results certainly do not prove that the bounded-degree sparsest multiple problem is NP-complete, they support our conjecture that it is.

Theorem 2.5. *The problem* SparseLatticeVector *of computing the vector with the least Hamming weight in an integer lattice specified by its basis is* NP*-complete.*

Proof. To see that the problem is in NP, a nondeterministic machine can just guess the positions at which the lattice vector is nonzero. The rest is a standard linear algebra problem.

We now show NP-hardness by giving a Cook-reduction from the problem *Subset Sum*, a well-known NP-complete problem.

We note first the standard formulation of *Subset Sum*: Given distinct integers $\{z_1, \ldots, z_n\}$, a target integer t and a positive integer $w \leq n$, is there a non-empty subset $S \subseteq \{1, \ldots, n\}$ of size exactly w such that such that $\sum_{i \in S} z_i = t$?

If $w = n$, the problem can be solved by comparing the sum $\sum_i z_i$ with t. Therefore, we can assume that $w < n$. Given an instance $\{z_1, \ldots, z_n\}$ of subset sum, to check if there is a subset of size $w < n$ summing to t, the reduction first creates the following matrix:

$$M_w = \begin{bmatrix} 1 & 1 & \cdots & 1 & 0 \\ z_1 & z_2 & \cdots & z_n & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ z_1^{w-1} & z_2^{w-1} & \cdots & z_n^{w-1} & 1 \\ z_1^w & z_2^w & \cdots & z_n^w & t \end{bmatrix} \in \mathbb{Z}^{(w+1)\times(n+1)}. \tag{2.4}$$

Lemma 2.6 (stated and proved below) shows that M_w has a null vector of sparsity at most $w + 1$ if and only if $z_{i_1} + z_{i_2} + \cdots + z_{i_w} = t$ for some $i_1 < i_2 < \ldots < i_w$.

To create an instance of *SparseLatticeVector*, the reduction creates a matrix N such that the columns of N span the kernel of M via $\mathbb{Z}$-linear combinations

(see Lemma 2.3). The instance $(\mathcal{L}, w)$, where $\mathcal{L}$ is the column lattice $\mathcal{L}$ of N, is fed to an algorithm claiming to solve the *Sparse Vector Problem.* □

Lemma 2.6. *The matrix M_w from equation (2.4) has a null vector of Hamming weight $w+1$ if and only if $z_{i_1} + z_{i_2} + \cdots + z_{i_w} = t$ for some $i_1 < i_2 < \ldots < i_w$.*

Proof. We will first prove that the sparsest null vector has weight at least $(w+1)$. To see this, consider the submatrix formed by any set of w columns. (We can assume that the last column is included in this set since otherwise the submatrix has a Vandermonde minor of size $w \times w$, and hence the columns are independent.) Since the principal minor of such a submatrix is a $(w-1)\times(w-1)$-sized Vandermonde matrix, the rows are independent. On adding either of the last two rows, the row-rank only increases since the other rows do not contain a nonzero entry in the last coordinate. Hence the row-rank (and hence the column-rank) of this submatrix is at least w, and hence the sparsest null vector of M_w has weight at least $(w+1)$.

Consider a $(w+1)$-sized subset of columns. If the last column is not in this set, the chosen columns form a Vandermonde matrix with nonzero determinant (since z_i are distinct). Therefore assume that the last column is among those chosen, the determinant of the resulting matrix can be expanded as:

$$\begin{vmatrix} 1 & \cdots & 1 & 0 \\ z_{i_1} & \cdots & z_{i_w} & 0 \\ \vdots & & \vdots & \vdots \\ z_{i_1}^{w-1} & \cdots & z_{i_w}^{w-1} & 1 \\ z_{i_1}^{w} & \cdots & z_{i_w}^{w} & t \end{vmatrix} = t \begin{vmatrix} 1 & \cdots & 1 \\ z_{i_1} & \cdots & z_{i_w} \\ \vdots & \vdots & \vdots \\ z_{i_1}^{w-1} & \cdots & z_{i_w}^{w-1} \end{vmatrix} - \begin{vmatrix} 1 & \cdots & 1 \\ z_{i_1} & \cdots & z_{i_w} \\ \vdots & \vdots & \vdots \\ z_{i_1}^{w-2} & \cdots & z_{i_w}^{w-2} \\ z_{i_1}^{w} & \cdots & z_{i_w}^{w} \end{vmatrix}.$$

The first of the matrices on the right-hand side is a Vandermonde matrix, whose determinant is well-known to be $\prod_{i_j<i_k}(z_{i_k} - z_{i_j})$. The second matrix is a first-order alternant whose determinant is known to be $(z_{i_1} + z_{i_2} + \cdots + z_{i_w}) \prod_{i_j<i_k}(z_{i_k} - z_{i_j})$. Hence the determinant of the entire matrix is $(t - z_{i_1} - z_{i_2} - \cdots - z_{i_w}) \prod_{i_j<i_k}(z_{i_k} - z_{i_j})$. Since all the z_i are distinct, the determinant vanishes if and only if the first term vanishes which holds when there exists a subset of $\{z_1, z_2, \ldots, z_n\}$ of size w summing to t. □

3 Binomial multiples over $\mathbb{Q}$

In this section we completely solve the problem of determining if there exists a binomial multiple of a rational input polynomial (i.e., a multiple of sparsity $t = 2$). That is, given input $f \in \mathbb{Q}[x]$ of degree d, we determine if there exists a binomial multiple $h = x^m - a \in \mathbb{Q}[x]$ of f, and if so, find such an h with minimal degree. The constant coefficient a will be given as a pair $(r, e) \in \mathbb{Q} \times \mathbb{N}$ representing $r^e \in \mathbb{Q}$. The algorithm requires a number of bit operations which is polynomial in d and $\log \mathcal{H}(f)$. No a priori bounds on the degree or height of h are required. We show that m may be exponential in d, and $\log a$ may be exponential in $\log \mathcal{H}(f)$, and give a family of polynomials with these properties.

Algorithm 3.1: Lowest degree Binomial Multiple of a Rational Polynomial

Input: $f \in \mathbb{Q}[x]$
Output: The lowest degree binomial multiple $h \in \mathbb{Q}[x]$ of f, or "`NONE`"

Factor f into irreducible factors: $f = x^b f_1 f_2 \cdots f_u$
if *f is not squarefree* **then return** "`NONE`"
for $i = 1, 2, 3, \ldots, u$ **do**
 $d_i \leftarrow \deg f_i$
 $m_i \leftarrow$ least $k \in \{d_i, d_i + 1, \ldots, d_i \cdot (\lceil 3 d_i \ln \ln d_i \rceil + 7)\}$ such that $x^k \operatorname{rem} f_i \in \mathbb{Q}$
 if *no such m_i is found* **then return** "`NONE`"
 else $r_i \leftarrow x^{m_i} \operatorname{rem} f_i$
$m \leftarrow \operatorname{lcm}(m_1, \ldots, m_u)$
foreach *2-subset* $\{i, j\} \subseteq \{1, \ldots, u\}$ **do**
 if $|r_i|^{m_j} \neq |r_j|^{m_i}$ **then return** "`NONE`"
 else if $\operatorname{sign}(r_i^{m/m_i}) \neq \operatorname{sign}(r_j^{m/m_j})$ **then** $m \leftarrow 2 \cdot \operatorname{lcm}(m_1, \ldots, m_u)$
return $x^b(x^m - r_1^{m/m_1})$, *with r_1 and m/m_1 given separately*

Algorithm 3.1 begins by factoring the given polynomial $f \in \mathbb{Q}[x]$ into irreducible factors (using, e.g., the algorithm of Lenstra et al. (1982)). We then show how to find a binomial multiple of each irreducible factor, and finally provide a combining strategy for the different multiples.

The following theorem of Risman (1976) characterizes binomial multiples of irreducible polynomials. Let $\phi(n)$ be Euler's totient function, the number of positive integers less than or equal to n which are coprime to n.

Fact 3.1 (Risman (1976), Proposition 4, Corollary 2.2)**.** Let $f \in \mathbb{Q}[x]$ be irreducible of degree d. Suppose the least-degree binomial multiple of f (if one exists) is of degree m. Then there exist $n, t \in \mathbb{N}$ with $n \mid d$ and $\phi(t) \mid d$ such that $m = n \cdot t$.

The following, easily derived from explicit bounds in Rosser and Schoenfeld (1962), gives a polynomial bound on m.

Lemma 3.2. *For all integers $n \geq 2$, $\phi(\lceil 3n \ln \ln n \rceil + 7) > n$.*

Proof. Rosser and Schoenfeld (1962), Theorem 15, implies that for all $n \geq 3$

$$\phi(n) > \frac{0.56146 \cdot n}{\ln \ln n + 1.40722}.$$

It is then easily derived by basic calculus that

$$\phi(3n \log \log n) > \frac{0.56146 \cdot (3n \log \log n)}{\ln \ln (3n \log \log n) + 1.40722} > n$$

for $n \geq 24348$. The inequality in the lemma statement is verified mechanically (say using Maple) for $2 \leq n \leq 24348$. □

Combining Fact 3.1 with Lemma 3.2, we obtain the following explicit upper bound on the maximum degree of a binomial multiple of an irreducible polynomial.

Theorem 3.3. *Let $f \in \mathbb{Q}[x]$ be irreducible of degree d. If a binomial multiple of f exists, and has minimal degree m, then $m \leq d \cdot (\lceil 3d \ln \ln d \rceil + 7)$.*

Proof. By Fact 3.1, $m = n \cdot t$ such that $n \mid d$ and $\phi(t) \mid d$. Define $\xi(n) = \lceil 3n \ln \ln n \rceil + 7$, and define $\xi^{-1}(n)$ to be the smallest integer such that $\xi(\xi^{-1}(n)) \geq n$. From Lemma 3.2, we have that $\phi(\xi(n)) > n$ for $n \geq 2$. Hence, $d \geq \phi(t) \geq \xi^{-1}(t)$. Since ξ is a non-decreasing function, $d \geq \xi^{-1}(t)$ implies that $\xi(d) \geq t$. Thus $m = n \cdot t \leq d \cdot \xi(d) \leq d \cdot (\lceil 3d \ln \ln d \rceil + 7)$. □

The above theorem ensures that for an irreducible f_i, Step 5 of Algorithm 3.1 computes the least-degree binomial multiple $x^{m_i} - r_i$ if it exists, and otherwise correctly reports failure. It clearly runs in polynomial time.

If f has any repeated factor, then it cannot have a binomial multiple (see Lemma 4.1 below). So assume the factorization of f is as computed in Step 1, and moreover f is squarefree. If any factor does not have a binomial multiple, neither can the product. If every irreducible factor does have a binomial multiple, Step 5 computes the one with the least degree. The following relates the degree of the minimal binomial multiple of the input polynomial to those of its irreducible factors.

Lemma 3.4. *Let $f \in \mathbb{Q}[x]$ be such that $f = f_1 \cdots f_u \in \mathbb{Q}[x]$ for distinct, irreducible $f_1, \ldots, f_u \in \mathbb{Q}[x]$. Let $f_i \mid (x^{m_i} - r_i)$ for minimal $m_i \in \mathbb{N}$ and $r_i \in \mathbb{Q}$, and let $f \mid (x^m - r)$ for $r \in \mathbb{Q}$. Then $\operatorname{lcm}(m_1, \ldots, m_u) \mid m$.*

Proof. It suffices to prove that if $f \mid (x^m - r)$ and $f_i \mid (x^{m_i} - r_i)$ for minimal m_i then $m_i \mid m$ since any multiple of f is also a multiple of f_i.

Assume for the sake of contradiction that $m = cm_i + \ell$ for $0 < \ell < m_i$. Then for any root $\alpha_i \in \mathbb{C}$ of f_i, we have that $r = \alpha^m = \alpha^{cm_i} \cdot \alpha^\ell = r_i^c \cdot \alpha^\ell$. Since r and r_i are both rational, so is α^ℓ. Also $\alpha^\ell = \beta^\ell$ for any two roots $\alpha, \beta \in \mathbb{C}$ of f_i. Hence $f_i \mid x^\ell - \alpha^\ell$ and $\ell < m_i$, contradicting the minimality of m_i.

Thus $m_i \mid m$, and therefore $\operatorname{lcm}(m_1, \ldots, m_u) \mid m$. □

Lemma 3.5. *For a polynomial $f \in \mathbb{Q}[x]$ factored into distinct irreducible factors $f = f_1 f_2 \ldots f_u$, with $f_i \mid (x^{m_i} - r_i)$ for $r_i \in \mathbb{Q}$ and minimal such m_i, a binomial multiple of f exists if and only if $|r_i|^{m_j} = |r_j|^{m_i}$ for every pair $1 \leq i, j \leq u$. If a binomial multiple exists, the least-degree binomial multiple of f is $x^m - r_i^{m/m_i}$ such that m either equals the least common multiple of the m_i or twice that number. It can be efficiently checked which of these cases holds.*

Proof. Let $\alpha_i \in \mathbb{C}$ be a root of f_i. For any candidate binomial multiple $x^m - r$ of f, we have (from Lemma 3.4) that $m_i \mid m$.

First, suppose that such a binomial multiple exists: $f \mid (x^m - r)$ with $r \in \mathbb{Q}$. It is easily seen from $\alpha_i^m = r$ and $\alpha_i^{m_i} = r_i$ that $r_i^{m/m_i} = r$. Since this holds

for any f_i, we see that $r_i^{m/m_i} = r = r_j^{m/m_j}$ for any $1 \le i, j \le u$. Thus $|r_i|^{m_j} = |r_j|^{m_i}$ must hold.

Conversely, suppose that $|r_i|^{m_j} = |r_j|^{m_i}$ holds for every pair $1 \le i, j \le u$. We get that $|\alpha_i|^{\ell m_i m_j} = |\alpha_j|^{\ell m_j m_i}$, and hence $\left|\alpha_i^\ell\right| = \left|\alpha_j^\ell\right|$ for $\ell = \operatorname{lcm}(m_1, \ldots, m_u)$. But α_i^ℓ are all rational since $m_i \mid \ell$. Thus $\alpha_i^{2\ell} = \alpha_j^{2\ell}$ for every pair i, j . Thus, there exists a binomial multiple of the original polynomial of degree 2ℓ.

To check whether $\alpha_i^\ell = \alpha_j^\ell$ holds (or in other words if the degree of the binomial multiple is actually the lcm), it suffices to check whether the sign of each α_i^ℓ is the same. This is equivalent to checking whether the sign of each r_i^{ℓ/m_i} is the same. Since we can explicitly compute ℓ and all the r_i, the sign of each r_i^{ℓ/m_i} can be easily computed from the sign of r_i and the parity of ℓ/m_i. □

The following comes directly from the previous lemma and the fact that Algorithm 3.1 performs polynomially many arithmetic operations.

Theorem 3.6. *Given a polynomial $f \in \mathbb{Q}[x]$, Algorithm 3.1 outputs the least-degree binomial multiple $x^m - r_i^{m/m_i}$ (with r_i and m/m_i output separately) if one exists or correctly reports the lack of a binomial multiple otherwise. Furthermore, it runs in deterministic time $(d + \mathcal{H}(f))^{O(1)}$.*

The constant coefficient of the binomial multiple cannot be output in standard form, but must remain an unevaluated power; the next theorem exhibits an infinite family of polynomials whose minimal binomial multiples have exponentially sized degrees and heights.

Theorem 3.7. *For any $d \ge 841$ there exists a polynomial $f \in \mathbb{Z}[x]$ of degree at most $d \log d$ and height $\mathcal{H}(f) \le \exp(2d \log d)$ whose minimal binomial multiple $x^m - a$ is such that $m > \exp(\sqrt{d})$ and $\mathcal{H}(a) > 2^{\exp(\sqrt{d})}$.*

Proof. We construct the family from a product of cyclotomic polynomials. Let $p_i \in \mathbb{N}$ be the i^{th} largest prime, and let $\Phi_{p_i} = (x^{p_i} - 1)/(x - 1) \in \mathbb{Z}[x]$ be the ${p_i}^{\text{th}}$ cyclotomic polynomials (whose roots are the primitive ${p_i}^{\text{th}}$ roots of unity). This is well known to be irreducible in $\mathbb{Q}[x]$.

Let $\ell = \sqrt{2d}$ and $g = \prod_{1 \le i \le \ell} \Phi_{p_i}$. Then, using the fact easily derived from Rosser and Schoenfeld (1962), Theorem 3, that $i \log i < p_i < 1.25 i \log i$ for all $i \ge 25$ and verifying that $(p_i - 1) \le 1.5 i \log i$ mechanically for smaller values of i,

$$\deg g = \sum_{1 \le i \le \ell} (p_i - 1) \ge \sum_{1 \le i \le \ell} i = \frac{l(l+1)}{2} \ge d,$$

and

$$\deg g = \sum_{1 \le i \le \ell} (p_i - 1) \le \sum_{1 \le i \le \ell} 1.5 i \log i \le 1.5 \left(\frac{\ell^2 + \ell}{2} \log \ell \right) \le d \log d.$$

The degree m of the minimal binomial multiple is the lcm of the order of the roots, and hence equal to the product of primes less than or equal to p_ℓ. This is $\exp(\vartheta(p_\ell))$ (where ϑ is the Chebyshev theta function), and for $\ell \geq 41$

$$m \geq \exp(\vartheta(p_\ell)) \geq \exp(\vartheta(\ell)) \geq \exp\left(\ell\left(1 - \frac{1}{\log \ell}\right)\right) \geq \exp\left(\sqrt{d}\right),$$

for $d \geq 841$, where the bounds on ϑ are derived from Rosser and Schoenfeld (1962) Theorem 4.

Now let $f = g(2x)$, so the minimal binomial multiple of f is $x^m - 1/2^m$. We have that

$$\mathcal{H}(g) \leq \prod_{1 \leq i \leq \ell} (1 + p_i) \leq 2^\ell \prod_{1 \leq i \leq \ell} p_i \leq \exp(2\ell \log \ell)$$

and

$$\mathcal{H}(f) \leq 2^{\deg(g)} \mathcal{H}(g) \leq 2^{d \log d} \exp(d \log d + 2\sqrt{2d} \log \sqrt{2d}) \leq \exp(2d \log d)$$

for all ≥ 841. □

4 Computing t-sparse multiples over $\mathbb{Q}$

We examine the problem of computing t-sparse multiples of rational polynomials, for any fixed positive integer t. As with other types of polynomial computations, it seems that cyclotomic polynomials behave quite differently from cyclotomic-free ones. Accordingly, we first examine the case that our input polynomial f consists only of cyclotomic or cyclotomic-free factors. Then we see how to combine them, in the case that none of the cyclotomic factors are repeated.

Specifically, we will show that, given any rational polynomial f which does not have repeated cyclotomic factors, and a height bound $c \in \mathbb{N}$, we can compute a sparsest multiple of f with height at most c, or conclude that none exists, in time polynomial in the size of f and $\log c$ (but exponential in t).

First, notice that multiplying a polynomial by a power of x does not affect the sparsity, and so without loss of generality we may assume all polynomials are relatively prime to x; we call such polynomials *non-original* since they do not pass through the origin.

4.1 The cyclotomic case

Suppose the input polynomial f is a product of cyclotomic factors, and write the complete factorization of f as

$$f = \Phi_{i_1}^{e_i} \cdot \Phi_{i_2}^{e_2} \cdots \Phi_{i_k}^{e_k}, \tag{4.1}$$

where Φ_j indicates the j^{th} cyclotomic polynomial, the i_j's are all distinct, and the e_i's are positive integers.

Now let $m = \operatorname{lcm}(i_1, \ldots, i_k)$. Then m is the least integer such that $\Phi_{i_1} \cdots \Phi_{i_k}$ divides $x^m - 1$. Let $\ell = \max_i e_i$, the maximum multiplicity of any factor of f. This means that $(x^m - 1)^\ell$ is an $(\ell+1)$-sparse multiple of f. To prove that this is in fact a sparsest multiple of f, we first require the following simple lemma. Here and for the remainder, for a univariate polynomial $f \in \mathsf{F}[x]$, we denote by f' the first derivative with respect to x, that is, $\frac{\mathrm{d}}{\mathrm{d}x} f$.

Lemma 4.1. *Let $h \in \mathbb{Q}[x]$ be a t-sparse and non-original polynomial, and write $h = a_1 + a_2 x^{d_2} + \cdots + a_t x^{d_t}$. Assume the complete factorization of h over $\mathbb{Q}[x]$ is $h = a_t h_1^{e_1} \cdots h_k^{e_k}$, with each h_i monic and irreducible. Then $\max_i e_i \leq t - 1$.*

Proof. Without loss of generality, assume h is *exactly* t-sparse, and each $a_i \neq 0$.

The proof is by induction on t. If $t = 1$ then $h = a_1$ is a constant, so $\max_i e_i = 0$ and the statement holds. Otherwise, assume the statement holds for $(t-1)$-sparse polynomials.

Write the so-called "sparse derivative" $\tilde{h}$ of h as

$$\tilde{h} = \frac{h'}{x^{d_2 - 1}} = a_2 d_2 + a_3 d_3 x^{d_3 - d_2} + \cdots + a_{t-1} d_{t-1} x^{d_{t-1} - d_2}.$$

For any i with $e_i > 0$, we know that $h_i^{e_i - 1}$ divides $\frac{d}{dx} h$, and h_i is relatively prime to $x^{d_2 - 1}$ since the constant coefficient of h is nonzero. Therefore $h_i^{e_i - 1}$ divides $\tilde{h}$. By the inductive hypothesis, since $\tilde{h}$ is $(t-1)$-sparse and non-original, $e_i - 1 \leq t - 2$, and therefore $e_i \leq t - 1$. Since i was chosen arbitrarily, $\max_i e_i \leq t - 1$. □

An immediate consequence is the following:

Corollary 4.2. *Let $f \in \mathbb{Q}[x]$ be a product of cyclotomic polynomials, written as in (4.1). Then*

$$h = (x^{\operatorname{lcm}(i_1, \ldots, i_k)} - 1)^{\max_i e_i}$$

is a sparsest multiple of f.

Proof. Clearly h is a multiple of f with exactly $\max_i e_i + 1$ nonzero terms. By way of contradiction, suppose a $(\max_i e_i)$-sparse multiple of f exists; call it $\bar{h}$. Without loss of generality, we can assume that $\bar{h}$ is non-original. Then from Lemma 4.1, the maximum multiplicity of any factor of $\bar{h}$ is $\max_i e_i - 1$. But this contradicts the fact that each $\Phi_i^{e_i}$ must divide $\bar{h}$. Therefore the original statement is false, and every multiple of f has at least $\max_i e_i + 1$ nonzero terms. □

4.2 The cyclotomic-free case

We say a polynomial $f \in \mathbb{Q}[x]$ is *cyclotomic-free* if it contains no cyclotomic factors. Here we will show that a sparsest multiple of a cyclotomic-free polynomial must have degree bounded by a polynomial in the size of the input and output.

First we need the following elementary lemma.

Lemma 4.3. *Suppose $f, h \in \mathbb{Q}[x]$ with f irreducible, and k is a positive integer. Then $f^k | h$ if and only if $f | h$ and $f^{k-1} | h'$.*

Proof. The $\Rightarrow$ direction is straightforward.

For the $\Leftarrow$ direction, suppose $f | h$ and $f^{k-1} | h'$. Let ℓ be the maximum multiplicity of f in h, and write $h = f^\ell g$ with $g \in \mathbb{Q}[x]$ relatively prime to f.

We can write $h' = f^{\ell-1}(fg' + \ell f' g)$. Now, by way of contradiction, assume that $k > \ell$. Then f divides $fg' + \ell f' g$, and therefore f divides $\ell f' g$. But this is impossible from the assumption that f is irreducible and relatively prime to g. Therefore $k \leq \ell$, and $f^k | f^\ell | h$. □

The following technical lemma provides the basis for our degree bound on the sparsest multiple of a non-cyclotomic polynomial.

Lemma 4.4. *Let $f, h_1, h_2, \ldots, h_\ell \in \mathbb{Q}[x]$ be non-original polynomials, where f is irreducible and non-cyclotomic with degree d, and each h_i satisfies $\deg h_i \leq u$ and $\mathcal{H}(h_i) \leq c$. Also let $k, m_1, m_2, \ldots, m_\ell$ be positive integers such that*

$$f^k | (h_1 x^{m_1} + h_2 x^{m_2} + \cdots + h_\ell x^{m_\ell}).$$

Then f^k divides each h_i whenever every "gap length", for $1 \leq i < \ell$, satisfies

$$m_{i+1} - m_i - \deg h_i \geq \frac{1}{2} d \cdot \ln^3(3d) \cdot \ln\left(u^{k-1} c\,(t-1)\right). \tag{4.2}$$

Proof. The proof is by induction on k. For the base case, let $k = 1$. Then we have a separate, inner induction on ℓ. The inner base case, when $k = \ell = 1$, is clear since f is non-original. Now assume the lemma holds whenever $k = 1$ and $1 \leq \ell - 1 < r$ for some $r \geq 2$. Let $g_1 = h_1 x^{m_1}$ and $g_2 = h_2 + \cdots + h_\ell x^{m_r - m_2}$, so that $f \,|\, (g_1 + g_2 x^{m_2})$. Since

$$m_2 - \deg g_1 \geq \frac{1}{2} d \cdot \ln^3(3d) \cdot \ln(c(t-1)),$$

we can apply (Lenstra, 1999, Proposition 2.3) to conclude that $f \,|\, g_1$ and $f \,|\, g_2$. This means $f \,|\, h_1$ and, by the inner induction hypothesis, $f \,|\, h_i$ for $2 \leq i \leq \ell$ as well. Therefore the lemma holds whenever $k = 1$.

Now assume the lemma holds whenever $\ell \geq 1$ and $1 \leq k < s$, for some $s \geq 2$. Next let ℓ be arbitrary and $k = s$. So we write $f^s | (h_1 x^{m_1} + \cdots + h_\ell x^{m_\ell})$.

The derivative of the right hand side is

$$h_1' x^{m_1} + m_1 h_1 x^{m_1 - 1} + \cdots + h_\ell' x^{m_\ell} + m_\ell h_\ell x^{m_\ell - 1},$$

which must be divisible by f^{s-1}. But by the induction hypothesis, f^{s-1} also divides each h_i, so we can remove all terms with h_i from the previous formula and conclude that $f^{s-1} \,|\, (h_1' x^{m_1} + \cdots + h_\ell' x^{m_\ell})$.

Since each $\mathcal{H}(h_i) \leq c$ and $\deg h_i \leq u$, the height of the derivative satisfies $\mathcal{H}(h_i') \leq uc$. A second application of the induction hypothesis therefore shows that each h_i' is divisible by f^{s-1}. Since $s - 1 \geq 1$, we already know that each h_i is divisible by f, and then applying Lemma 4.3 completes the proof. □

Our main tool in proving that Algorithm 2.2 is useful for computing the sparsest multiple of a rational polynomial, given only a bound c on the height, in polynomial time in the size of f and $\log c$, is the following degree bound on the sparsest height-bounded multiple of a rational polynomial.

Theorem 4.5. *Let $f \in \mathbb{Q}[x]$ with $\deg f = d$ be cyclotomic-free, and let $t, c \in \mathbb{N}$ such that f has a nonzero t-sparse multiple with height at most c. Denote by n the smallest degree of any such multiple of f. Then n satisfies*

$$n \leq 2(t-1)B \ln B, \tag{4.3}$$

where B is the formula polynomially bounded by d, $\log c$, and $\log t$ defined as

$$B = \frac{1}{2}d^2 \cdot \ln^3(3d) \cdot \ln\left(\hat{c}\,(t-1)^d\right), \tag{4.4}$$

and $\hat{c} = \max(c, 35)$.

Proof. Let h be a t-sparse multiple of f with degree n and height $\mathcal{H}(h) \leq c$. Without loss of generality, assume $d \geq 1$, $t \geq 2$, and both f and h are non-original.

By way of contradiction, assume $n > 2(t-1)B \ln B$. For any univariate polynomial define the *gap lengths* to be the differences of consecutive exponents of nonzero terms. Split h at every gap greater than $2B \ln B$ by writing

$$h = h_1 x^{m_1} + h_2 x^{m_2} + \cdots + h_\ell x^{m_\ell},$$

where each $h_i \in \mathbb{Q}[x]$ has nonzero constant term and each gap length satisfies $m_{i+1} - m_i - \deg h_i > 2B \ln B$. Since we split h at *every* sufficiently large gap, and h has at most t nonzero terms, each h_i has degree at most $u = 2(t-1)B \ln B$.

We want to show that the gap length $2B \ln B$ is sufficiently large to apply Lemma 4.4. For this, first notice that $2B \ln B = B \ln(B^2)$. Since B is positive, $B^2 > 2B \ln B$, so the gap length is greater than $B \ln(2B \ln B)$.

Since $\hat{c} \geq 35$, $B \geq 2.357$, and then

$$\begin{aligned}(d-1)\ln(2B\ln B) \cdot \ln(\hat{c}(t-1)^d) &> \ln\left((2B\ln B)^{d-1} \cdot \hat{c}(t-1)^d\right) \\ &= \ln\left(u^{d-1}\hat{c}\,(t-1)\right).\end{aligned}$$

Then from the definition of B in (4.4), the gap length satisfies

$$2B \ln B > B \ln(2B \ln B) > \frac{1}{2}d \cdot \ln^3(3d) \cdot \ln\left(u^{d-1}\hat{c}\,(t-1)\right).$$

Finally, notice that the maximum multiplicity of any factor of f is at most $\deg f = d$. Thus, using the notation of Lemma 4.4, $d \geq k$. Therefore Lemma 4.4 applies to each factor of f (to full multiplicity) and we conclude that f divides each h_i.

But then, since there is at least one gap and $\ell > 1$, h_1 is a multiple of f with fewer terms and lower degree than h. This is a contradiction, which completes the proof. □

In order to compute the sparsest multiple of a rational polynomial with no cyclotomic or repeated factors, we therefore can simply call Algorithm 2.2 with the given height bound c and degree bound as specified in (4.3).

4.3 Handling cyclotomic factors

Suppose f is any non-original rational polynomial with no repeated cyclotomic factors. Factor f as $f = f_C \cdot f_D$, where f_C is a squarefree product of cyclotomics and f_D is cyclotomic-free. Write the factorization of f_C as $f_C = \Phi_{i_1} \cdots \Phi_{i_k}$, where Φ_n is the n^{th} cyclotomic polynomial. Since every i^{th} root of unity is also a $(mi)^{\text{th}}$ root of unity for any $m \in \mathbb{N}$, f_C must divide the binomial $x^{\text{lcm}\{i_1,\ldots,i_k\}} - 1$, which is in fact a sparsest multiple of f_C (Corollary 4.2) and clearly has minimal height.

Then we will show that a sparsest height-bounded multiple of f is either of small degree, or can be constructed as a sparsest height-bounded multiple of f_D times the binomial multiple of f_C specified above. Algorithm 4.1 uses this fact to compute a sparsest multiple of any such f.

Algorithm 4.1: Rational Sparsest Multiple

Input: Bounds $t, c \in \mathbb{N}$ and $f \in \mathbb{Q}[x]$ a non-original polynomial of degree d with no repeated cyclotomic factors

Output: t-sparse multiple h of f with $\mathcal{H}(h) \leq c$, or "NONE"

1 Factor f as $f = \Phi_{i_1} \cdot \Phi_{i_2} \cdots \Phi_{i_k} \cdot f_D$, where f_D is cyclotomic-free

2 $n \leftarrow$ degree bound from (4.3)

3 $\hat{h} \leftarrow \lfloor t/2 \rfloor$-sparse multiple of f_D with $\mathcal{H}(\hat{h}) \leq c$ and $\deg \hat{h} \leq n$, using Algorithm 2.2

4 $\tilde{h} \leftarrow t$-sparse multiple of f with $\mathcal{H}(h) \leq c$ and $\deg h \leq n$, using Algorithm 2.2

5 **if** $\hat{h} =$"NONE" *and* $\tilde{h} =$"NONE" **then return** "NONE"

6 **else if** $\hat{h} =$"NONE" *or* $\text{sparsity}(\tilde{h}) \leq 2 \cdot \text{sparsity}(\hat{h})$ **then return** $\tilde{h}$

7 $m \leftarrow \text{lcm}\{i_1, i_2, \ldots, i_k\}$

8 **return** $\hat{h} \cdot (x^m - 1)$

Theorem 4.6. *Let $f \in \mathbb{Q}[x]$ be a degree-d non-original polynomial with no repeated cyclotomic factors. Given f and integers c and t, Algorithm 4.1 correctly computes a t-sparse multiple h of f satisfying $\mathcal{H}(h) \leq c$, if one exists. The sparsity of h will be minimal over all multiples with height at most c. The algorithm requires $(d \log c)^{O(t)} \cdot 2^{O(t \log t)} \cdot (\log \mathcal{H}(f))^{O(1)}$ bit operations.*

Proof. Step 1 can be accomplished in the stated complexity bound using Lenstra et al. (1982). The cost of the remaining steps follows from basic arithmetic and Theorem 2.4. Define h to be sparsest multiple of f of least degree that satisfies $\mathcal{H}(h) \leq c$. We have two cases:

Case 1: $\deg h \leq n$**.** Then the computed $\tilde{h}$ must equal h. Furthermore, since this is the sparsest multiple, either $\hat{h}$ does not exist or the sparsity of $\hat{h}$ is greater than or equal to the sparsity of $\tilde{h}$. So $h = \tilde{h}$ is correctly returned by the algorithm in this case.

Case 2: $\deg h > n$**.** Then, using Lemma 4.4, since $f_D \mid h$, h can be written $h = h_1 + x^i h_2$, for some $i > \deg h_1$, and f_D divides both h_1 and h_2. By Theorem 2.4, sparsity$(\hat{h})$ must then be less than or equal to each of sparsity(h_1) and sparsity(h_2). But since sparsity(h) = sparsity(h_1) + sparsity(h_2), this means that the sparsity of $\hat{h} \cdot (x^m - 1)$ is less than or equal to the sparsity of h, and hence this is a sparsest multiple.

□

4.4 An example

Say we want to find a sparsest multiple, with coefficients at most 1000 in absolute value, of the following polynomial over $\mathbb{Z}[x]$.

$$f = x^{10} - 5x^9 + 10x^8 - 8x^7 + 7x^6 - 4x^5 + 4x^4 + x^3 + x^2 - 2x + 4$$

Note that finding the *sparsest* multiple would correspond to setting $t = 10$ in the algorithm (since the least-degree 11-sparse multiple is f itself). To accomplish this, we first factor f using (Lenstra et al., 1982) and identify cyclotomic factors:

$$f = \underbrace{(x^2 - x + 1)}_{\Phi_6} \cdot \underbrace{(x^4 - x^3 + x^2 - x + 1)}_{\Phi_{10}} \cdot \underbrace{(x^4 - 3x^3 + x^2 + 6x + 4)}_{f_D}.$$

Next, we calculate a degree bound from Theorem 4.5. Unfortunately, this bound is not very tight (despite being polynomial in the output size); using $t = 10$, $c = 1000$, and f given above, the bound is $n \leq 11\,195\,728$. So for this example, we will use the smaller (but artificial) bound of $n \leq 20$.

The next step is to calculate the sparsest 5-sparse multiple of f_D and 10-sparse multiple of f with degrees at most 20 and heights at most 1000. Using Algorithm 2.2, these are respectively

$$\begin{aligned}\hat{h} &= x^{12} + 259x^6 + 64\\ \tilde{h} &= x^{11} - 3x^{10} + 12x^8 - 9x^7 + 10x^6 - 4x^5 + 9x^4 + 3x^3 + 8.\end{aligned}$$

Since the sparsity of $\hat{h}$ is less than half that of $\tilde{h}$, a sparsest multiple is

$$\begin{aligned}h &= (x^{12} + 259x^6 + 64) \cdot (x^{\mathrm{lcm}(6,10)} - 1)\\ &= x^{42} + 259x^{36} + 64x^{30} - x^{12} - 259x^6 - 64.\end{aligned}$$

5 Sparse multiples over $\mathbb{F}_q$

We prove that for any constant t, finding the minimal degree t-sparse multiple of an $f \in \mathbb{F}_q[x]$ is harder than finding orders of elements in $\mathbb{F}_{q^e}$. Order finding is reducible to integer factorization and to discrete logarithm, but reductions in the other direction are not known for finite fields (Adleman and McCurley, 1994). However, at least for prime fields and assuming the Extended Riemann Hypothesis, a fast algorithm for order finding in finite fields would give an efficient procedure for computing primitive elements (Wang, 1959; Shoup, 1992). The latter problem is regarded as "one of the most important unsolved and notoriously hard problems in the computational theory of finite fields" (von zur Gathen and Shparlinski, 1999).

Formal problem definitions are as follows:

$\mathbf{SpMul}^{(t)}_{\mathbb{F}_q}(f, n)$: Given a polynomial $f \in \mathbb{F}_q[x]$ and an integer $n \in \mathbb{N}$, determine if there exists a (nonzero) 2-sparse multiple $h \in \mathbb{F}_q[x]$ of f with $\deg h \leq n$.

$\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$: Given an element $a \in \mathbb{F}^*_{q^e}$ and an integer $n < q^e$, determine if there exists a positive integer $m \leq n$ such that $a^m = 1$.

The problem $\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$ is well-studied (see for instance Meijer (1996)), and has been used as a primitive in several cryptographic schemes. Note that an algorithm to solve $\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$ will allow us to determine the *multiplicative order* of any $a \in \mathbb{F}^*_{q^e}$ (the smallest nonzero m such that $a^m = 1$) with essentially the same cost (up to a factor of $O(e \log q)$) by using binary search.

The reduction from $\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$ to $\mathbf{SpMul}^{(t)}_{\mathbb{F}_q}(f, n)$ works as follows: Given an instance of $\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$, we first check if the order o_a of a is less than t by brute-force. Otherwise, we construct the minimal polynomial g_{a^i} (over $\mathbb{F}_q$) for each $a^0, a^1, a^2, \ldots, a^{t-1}$. We only keep distinct g_{a_i}, and call the product of these distinct polynomials $f_{a,t}$. We then run the $\mathbf{SpMul}^{(t)}_{\mathbb{F}_q}(f, n)$ subroutine to search for the existence of a degree n, t-sparse multiple of the polynomial $f_{a,t}$.

Theorem 5.1. *Let $a \in \mathbb{F}_q$ be an element of order at least t. Then the least degree t-sparse multiple of $f_{a,t}$ is $x^{o_a} - 1$ where o_a is the order of a.*

Proof. It is easy to see that $x^{o_a} - 1$ is a multiple of the given polynomial. We need to prove that it is actually the least-degree t-sparse multiple.

By equation (2.2) in Section 2, a degree n multiple h of $f_{a,t}$ corresponds to the following set of linear equations:

$$\underbrace{\begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & a & a^2 & \cdots & a^{n-1} \\ 1 & a^2 & a^4 & \cdots & a^{2n-2} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & a^t & a^{2t} & \cdots & a^{tn-t} \end{bmatrix}}_{A(f_{a,t},n)} \begin{bmatrix} h_0 \\ h_1 \\ \\ \\ \vdots \\ h_{n-1} \end{bmatrix} = 0.$$

To prove that no t-sparse multiple h of degree less than o_a exists, it suffices to show that any t columns of $A(f_{a,t}, o_a - 1)$ are linearly independent. Consider the $(t \times t)$-matrix corresponding to some choice of t columns:

$$B = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ a^{i_1} & a^{i_2} & \cdots & a^{i_t} \\ \vdots & \vdots & \vdots & \vdots \\ a^{ti_1} & a^{ti_2} & \cdots & a^{ti_t} \end{bmatrix}.$$

This Vandermonde matrix $\mathbf{B}$ has determinant $\prod_{1 \le j < k \le t}(a^{i_k} - a^{i_j})$ which is nonzero since $i_j < i_k < o_a$ and hence $a^{i_j} \neq a^{i_k}$. Thus the least-degree t-sparse multiple of the given polynomial is $x^{o_a} - 1$. □

Of cryptographic interest is the fact that the order-finding polynomials in the reduction above are sufficiently dense in $\mathbb{F}_q[x]$ that the reduction also holds in the average case. That is, an algorithm for sparsest multiples that is polynomial-time on average would imply an average case polynomial-time algorithm for order finding in $\mathbb{F}_{q^d}$.

Next we give a probabilistic algorithm for finding the least degree binomial multiple for polynomials $f \in \mathbb{F}_q$. This algorithm makes repeated calls to an $\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$ (defined in the previous section) subroutine. Combined with the hardness result of the previous section (with t=2), this characterizes the complexity of finding least-degree binomial multiples in terms of the complexity of $\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$, upto randomization.

Algorithm 5.1 solves the binomial multiple problem in $\mathbb{F}_q$ by making calls to an $\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$ procedure that computes the order of elements in extension fields of $\mathbb{F}_q$. Thus $\mathbf{SpMul}^{(2)}_{\mathbb{F}_q}(f)$ reduces to $\mathbf{Order}_{\mathbb{F}_{q^e}}(a, n)$ in probabilistic polynomial time. Construction of an irreducible polynomial (required for finite field arithmetic) as well as the factoring step in the algorithm make it probabilistic.

Theorem 5.2. *Given $f \in \mathbb{F}_q[x]$ of degree d, Algorithm 5.1 correctly computes a binomial multiple h of f with least degree. It uses at most d^2 calls to a routine for order finding in $\mathbb{F}_{q^e}$, for various $e \le d$, and $d^{O(1)}$ other operations in $\mathbb{F}_q$. It is probabilistic of the Las Vegas type.*

Proof. As a first step, the algorithm factors the given polynomial into irreducible factors. Efficient probabilistic algorithms for factoring polynomials over finite fields are well-known (von zur Gathen and Gerhard (2003)).

First, suppose the input polynomial f is irreducible, i.e. $\ell = e_1 = 1$ in Step 1. Then it has the form $f = (x - a)(x - a^q) \cdots (x - a^{q^{d-1}})$ for some $a \in \mathbb{F}_{q^d}$, where $d = \deg f$. If $f = (x - a)$, the least-degree binomial multiple is f itself. Therefore, assume that $d > 1$. Let the least-degree binomial multiple (in $\mathbb{F}_q[x]$) be $x^n - \beta$

Since both a and a^q are roots of $(x^n - \beta)$, we have that $a^n = a^{nq}$ and $a^{n(q-1)} = 1$. Thus, the order o_a of a divides $n(q-1)$. The minimal n for which

Algorithm 5.1: Least degree binomial multiple of f over $\mathbb{F}_q$

Input: $f \in \mathbb{F}_q[x]$
Output: The least degree binomial multiple h of f

1 Factor $f = x^b f_1^{e_1} \cdot f_2^{e_2} \cdot f_\ell^{e_\ell}$ for irreducible $f_1, \ldots, f_\ell \in \mathbb{F}_q[x]$, and set $d_i \leftarrow \deg f_i$
2 **for** $i = 1, 2, \ldots, \ell$ **do**
3 $\quad a_i \leftarrow x \in \mathbb{F}_q[x]/(f_i)$, a root of f_i in the extension $\mathbb{F}_{q^{d_i}}$
4 $\quad$ Calculate o_i, the order of a_i in $\mathbb{F}_q[x]/(f_i)$.
5 $n_1 \leftarrow \operatorname{lcm}(\{o_i / \gcd(o_i, q-1)\})$ for all i such that $d_i > 1$
6 $n_2 \leftarrow \operatorname{lcm}(\{order(a_i/a_j)\})$ over all $1 \le i, j \le u$
7 $n \leftarrow \operatorname{lcm}(n_1, n_2)$
8 $\tilde{h} \leftarrow (x^n - a_1^n)$
9 $e \leftarrow \lceil \log_p \max e_i \rceil$, the smallest e such that $p^e \ge e_i$ for all i
10 **return** $h = x^b (x^n - a_1^n)^{p^e}$

$o_a \mid n(q-1)$ is $n = \frac{o_a}{\gcd(o_a, q-1)}$. Since this n ensures that $a^n = a^{nq}$, it also simultaneously ensures that each a^{q^i} is also a root.

Notice that this n equals n_1 computed on Step 5, and n_2 computed on Step 6 will equal 1, so the algorithm is correct in this case.

Now suppose the input polynomial f is reducible. The factorization step factors f into irreducible factors $f = f_1^{e_1} f_2^{e_2} \cdots f_\ell^{e_\ell}$. Let $\check{f} = f_1 f_2 \cdots f_\ell$ denote the squarefree part of f.

Being irreducible, each f_i has the form $f_i(x) = (x - a_i)(x - a_i^q) \cdots (x - a_i^{q^{d_i - 1}})$ for some $a_i \in \mathbb{F}_{q^d}$, and $d_i = \deg f_i$. We make two observations:

- If $\check{f}(x) \mid x^n - a$ for some $a \in \mathbb{F}_q$, we have that $a_i^n = a_j^n$ for all $1 \le i, j \le \ell$, and hence that $(\frac{a_i}{a_j})^n = 1$. Thus $order(\frac{a_i}{a_j}) \mid n$. The least integer satisfying these constraints is n_2 computed on Step 6.

- As before for the case when the input polynomial is irreducible and of degree more than one: $d_i > 1$ implies that $\frac{o_i}{\gcd(o_i, q-1)} \mid n$ for o_i the order of a_i. The least integer satisfying these constraints is n_1 computed on Step 5.

The minimal n is the least common multiple of all the divisors obtained from the above two types of constraints, which is exactly the value computed on Step 7. The minimal degree binomial multiple of $\check{f}$ is $x^n - a_1^n$.

It is easily seen that for the smallest e such that $p^e \ge e_i$, $(x^n - a^n)^{p^e}$ is a binomial multiple of f. We now show that it is actually the minimal degree binomial multiple. Specifically, let e be the smallest non-negative integer such that $p^e \ge \max e_i$; we show that the minimal degree binomial multiple of f is $(x^n - a_i^n)^{p^e}$ for n obtained as above.

Let the minimal degree binomial multiple of f be $x^{\hat{n}} - b$. Factor $\hat{n}$ as $\hat{n} = \check{n} p^c$ for maximal c, and write $(x^{\hat{n}} - b)$ as $(x^{\check{n}} - b^{1/p^c})^{p^c}$. The squarefree part of

f, $\check{f}$ divides $(x^{\check{n}} - b^{1/p^c})$, and hence (by constraints on and minimality of n) $(x^n - a_1^n) \mid (x^{\check{n}} - b^{1/p^c})$. Thus $\check{n} \geq n$.

Since c is chosen maximally, p does not divide $\check{n}$, and hence $x^{\check{n}} - b^{1/p^c}$ is squarefree. Using this and the fact that f divides $(x^{\check{n}} - b^{1/p^c})^{p^c}$, it is seen that $p^c \geq e_i$ holds for all e_i, and hence $p^c \geq p^e$. This, along with $\check{n} \geq n$, completes the proof that $(x^n - a_i^n)^{p^e}$ is the minimal degree binomial multiple of f, which completes the proof of the theorem. □

6 Conclusion and Open Problems

To summarize, we have presented an efficient algorithm to compute the least-degree binomial multiple of any rational polynomial. We can also compute t-sparse multiples of rational polynomials that do not have repeated cyclotomic factors, for any fixed t, and given a bound on the height of the multiple.

We have also shown that, even for fixed t, finding a t-sparse multiple of a degree-d polynomial over $\mathbb{F}_q[x]$ is at least as hard as finding the orders of elements in $\mathbb{F}_{q^d}$. In the $t = 2$ case, there is also a probabilistic reduction in the other direction, so that computing binomial multiples of degree-d polynomials over $\mathbb{F}_q[x]$ probabilisticly reduces to order finding in $\mathbb{F}_{q^d}$.

Several important questions remain unanswered. Although we have an unconditional algorithm to compute binomial multiples of rational polynomials, computing t-sparse multiples for fixed $t \geq 3$ requires an a priori height bound on the output as well as the requirement that the input contains no repeated cyclotomic factors. Removing these restrictions is desirable (though not necessarily possible).

Regarding lower bounds, we know that computing t-sparse multiples over finite fields is at least as hard as order finding, a result which is tight (up to randomization) for $t = 2$, but for larger t we believe the problem is even harder. Specifically, we suspect that computing t-sparse multiples is NP-complete over both $\mathbb{Q}$ and $\mathbb{F}_q$, when t is a parameter in the input.

Acknowledgments

The authors would like to thank John May, Arne Storjohann, and the anonymous referees from ISAAC 2010 for their careful reading and useful observations on earlier versions of this work.

References

L. M. Adleman and K. S. McCurley. Open problems in number-theoretic complexity. II. In *Algorithmic number theory (Ithaca, NY, 1994)*, volume 877 of *Lecture Notes in Comput. Sci.*, pages 291–322. Springer, Berlin, 1994.

Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest

lattice vector problem. In *Symp. Theory of Computing (STOC'01)*, pages 601–610, 2001.

J.-P. Aumasson, M. Finiasz, W. Meier, and S. Vaudenay. TCHo: a hardware-oriented trapdoor cipher. In *ACISP'07: Proceedings of the 12th Australasian conference on Information security and privacy*, pages 184–199, Berlin, Heidelberg, 2007. Springer-Verlag. ISBN 978-3-540-73457-4.

E. R. Berlekamp, R. J. McEliece, and H. C. van Tilborg. On the inherent intractability of certain coding problems. *IEEE Transactions on Information Theory*, 24(3), 1978.

R. P. Brent and P. Zimmermann. Algorithms for finding almost irreducible and almost primitive trinomials. In *Primes and Misdemeanours: Lectures in Honour of the Sixtieth Birthday of Hugh Cowie Williams, Fields Institute*, page 212, 2003.

F. Didier and Y. Laigle-Chapuy. Finding low-weight polynomial multiples using discrete logarithms. In *Proc. IEEE International Symposium on Information Theory (ISIT 2007)*, pages 1036–1040, 2007.

Sebastian Egner and Torsten Minkwitz. Sparsification of rectangular matrices. *J. Symb. Comput.*, 26(2):135–149, 1998.

Laila El Aimani and Joachim von zur Gathen. Finding low weight polynomial multiples using lattices. Cryptology ePrint Archive, Report 2007/423, 2007. http://eprint.iacr.org/2007/423.pdf.

I. Z. Emiris and I. S. Kotsireas. Implicitization exploiting sparseness. In *Geometric and algorithmic aspects of computer-aided design and manufacturing*, volume 67 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 281–297, 2005.

Joachim von zur Gathen and J. Gerhard. *Modern Computer Algebra*, chapter 14, pages 367–380. Cambridge University Press, New York, NY, USA, 2003. ISBN 0521826462.

Joachim von zur Gathen and I. Shparlinski. Constructing elements of large order in finite fields. In M. Fossorier, H. Imai, S. Lin, and A. Poli, editors, *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*, volume 1719 of *Lecture Notes in Computer Science*, pages 730–730. Springer Berlin / Heidelberg, 1999.

Mark Giesbrecht, Daniel S. Roche, and Hrushikesh Tilak. Computing sparse multiples of polynomials. In *Proc. 21st International Symposium on Algorithms and Computation (ISAAC 2010)*, pages 266–278, 2010.

Venkatesan Guruswami and Alexander Vardy. Maximum-likelihood decoding of reed-solomon codes is NP-hard. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 470–478, 2005.

A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982.

H. W. Lenstra, Jr. Finding small degree factors of lacunary polynomials. In *Number theory in progress, Vol. 1 (Zakopane-Kościelisko, 1997)*, pages 267–276. de Gruyter, Berlin, 1999.

A. R. Meijer. Groups, factoring, and cryptography. *Math. Mag.*, 69(2):103–109, 1996.

Oded Regev. A simply exponential algorithm for svp (ajtai-kumar-sivakumar). online: http://www.cs.tau.ac.il/~odedr/teaching/lattices_fall_2004/, 2004. Scribe: Michael Khanevsky.

Lawrence J. Risman. On the order and degree of solutions to pure equations. *Proc. Amer. Math. Soc.*, 55(2):261–266, 1976.

J. B. Rosser and L. Schoenfeld. Approximate formulas for some functions of prime numbers. *Ill. J. Math.*, 6:64–94, 1962.

H.R. Sadjadpour, N.J.A. Sloane, M. Salehi, and G. Nebe. Interleaver design for turbo codes. *IEEE Journal on Selected Areas in Communications*, 19(5):831–837, 2001.

Victor Shoup. Searching for primitive roots in finite fields. *Math. Comp.*, 58 (197):369–380, 1992. doi: 10.2307/2153041.

Arne Storjohann. *Algorithms for Matrix Canonical Forms*. PhD thesis, Swiss Federal Institute of Technology Zürich, 2000.

Hrushikesh Tilak. Computing sparse multiples of polynomials. Master's thesis, University of Waterloo, 2010.

Alexander Vardy. The intractability of computing the minimum distance of a code. *IEEE Transactions on Information Theory*, 43(6):1757–1766, 1997.

Yuan Wang. On the least primitive root of a prime. *Acta Math. Sinica*, 9: 432–441, 1959.

A Finding short l_∞ vectors in lattices

In Section 2.1, we presented Algorithm 2.1 to find the shortest l_∞ vector in the image of an integer matrix. This appendix is devoted to proving the correctness of this algorithm, culminating in the proof of Theorem 2.1. Again, the results here are due to the presentation of Ajtai et al. (2001) by Regev (2004), with modifications to accommodate the infinity norm.

For any lattice $\mathcal{L}$, define $s(\mathcal{L}) = \min_{v \in \mathcal{L}} \|v\|_2$ to be the least l_2 norm of any vector in $\mathcal{L}$. If $\mathcal{L}$ satisfies $2 \leq s(\mathcal{L}) < 3$, and B is a basis for $\mathcal{L}$, then we will show that the main for loop in Steps 3–24 of Algorithm 2.1 finds a vector in $\mathcal{L}$ with minimal l_∞ norm, with high probability. The for loop on line 3 adapts this to work for any lattice by scaling. More precisely, given a lattice $\mathcal{L}$, we first run the algorithm of Lenstra et al. (1982) to get an approximation λ for the shortest l_2 vector in $\mathcal{L}$ satisfying $s(\mathcal{L}) \leq \|\lambda\|_2 \leq 2^n s(\mathcal{L})$. For each k from 1 to $2n$, we then run the for loop with basis B_k for the lattice $(1.5^k / \|\lambda\|_2) \cdot \mathcal{L}$. For

some k in this range, $2 \leq s(B_k) < 3$ must hold, and we will show that for this k, the vector v_k set on Step 24 is the l_∞ shortest vector in the image of B_k with high probability. For every k, v_k is a vector in the image of B_k, and hence it suffices to output the shortest l_∞ vector among $\{(\|\lambda\|_2 / 1.5^k) v_k\}$ on Step 25.

We will now prove that the vector v_k set on Step 24 is with high probability the shortest l_∞ vector in the image of B, when B is a basis for a lattice $\mathcal{L}$ such that $2 \leq s(\mathcal{L}) < 3$.

To find the shortest l_∞ vector in a lattice, it suffices to consider all lattice vectors of l_2 norm at most $\sqrt{n}$ times the norm of the shortest l_2 vector. Algorithm 2.1 achieves this by running the main body of the loop with different values of γ. In a particular iteration of the outermost loop, with high probability, the algorithm encounters all lattice vectors v with l_2 norm satisfying $(2/3) \cdot \|v\|_2 \leq \gamma < \|v\|_2$. Call all such v *interesting*. By iterating over a suitable range of γ, it returns the shortest l_∞ vector among all the *interesting* vectors, which with high probability is the shortest l_∞ vector in the lattice.

For a particular iteration of the loop (with a fixed γ), the algorithm uniformly samples a large number of vectors from an appropriately sized ball. In fact, the algorithm works even if an almost-uniform sampling over rational vectors with bit lengths bounded by $(\log \|B\| + n)^{O(1)}$ is performed. This is because the size of sufficiently small lattice vectors is only a polynomial in the size of the basis vectors. For the rest of this subsection, "arithmetic operations" means operations with rational numbers of this size.

After sampling, the algorithm performs a series of *sieving* steps to ensure that at the end of these steps the algorithm is left with lattice vectors of sufficiently small l_2 norm. Using a probabilistic argument, it is argued that all *interesting* vectors are obtained.

The following lemma proves the correctness of the sieving steps. These correspond to Steps 13 to 17 of the algorithm. At the end of this sieving, the algorithm produces a set J of size at most 5^n.

Lemma A.1. *Given $S \subseteq \{1, \ldots, m\}$ such that for all $i \in S$, $y_i \in \mathbb{R}^n$ and $\|y_i\|_2 \leq r$, Steps 13–17 efficiently compute the following: a subset $J \subseteq S$ of size at most 5^n and a mapping $\eta : S \setminus J \to J$ such that $\|y_i - y_{\eta_i}\|_2 \leq r/2$.*

Proof. Initially the set J is empty. The algorithm iterates over the points y_i with $i \in S$, adding i to J only if $min_{j \in J}(\|y_j - y_i\|_2) > r/2$. For $i \notin J$, it sets η_i to a $j \in J$ such that $\|y_j - y_i\|_2 \leq r/2$.

It is clear that this procedure runs in polynomial time. To see that the size of J is at most 5^n, note that all the balls of radius $R/4$ and centered at y_j for $j \in J$ are disjoint by construction of J. Also, these balls are contained in a ball of radius $R + R/4$ since $\|y_i\|_2 \leq R$. Thus the total number of disjoint balls, and hence the size of J, can be bounded above by comparing the volumes: $|J| \leq ((5R/4)/(R/4))^n = 5^n$. □

The algorithm views every sampled vector x_i as a perturbation of a lattice vector $x_i - y_i$ for some y_i. The idea is the following: initially y_i is calculated so that x_i is a perturbation of some large lattice vector. Iteratively, the algorithm

either obtains shorter and shorter lattice vectors corresponding to x_i, or discards x_i in some sieving step. At all stages of the algorithm, $x_i - y_i$ is a lattice vector. The following two lemmas concretize these observations.

Lemma A.2. *$\{y_i\}$ can be found efficiently in Step 11; and $\{x_i - y_i\} \subseteq \mathcal{L}$.*

Proof. For a fixed x_i, y_i is set to $(x_i \bmod \mathcal{P}(B))$ where $\mathcal{P}(B)$ denotes all vectors contained in the parallelogram $\left\{ \sum_{i=1}^{n} \alpha_i b_i \mid 0 \leq \alpha_i < 1 \right\}$, with b_i being the given basis vectors. Thus y_i is the unique element in $\mathcal{P}(B)$ such that $y_i = x_i - v$ for $v \in \mathcal{L}$. From this definition of y_i, we get that $x_i - y_i \in \mathcal{L}$ for every i.

To calculate y_i efficiently, simply represent x_i as a rational linear combination of the basis vectors $\{b_i\}$ and then truncate each coefficient modulo 1. □

Lemma A.3. *$Y_\gamma \subseteq \mathcal{L} \cap \mathbf{B}_n(0, 3\gamma + 1)$.*

Proof. By Lemma A.2, $(x_i - y_i) \in \mathcal{L}$ for all $i \in S$ before the start of the loop. It needs to be proved that the same holds after the loop, and furthermore, all the resulting lattice vectors lie in $\mathbf{B}_n(0, 3\gamma + 1)$. Whenever the algorithm modifies any y_i, it sets it to $y_i + x_{\eta(i)} - y_{\eta(i)}$; and thus a lattice vector $(x_i - y_i)$ changes into $(x_i - y_i) - (x_{\eta(i)} - y_{\eta(i)})$. Since both of the terms are lattice vectors, so is their difference. Thus $Y_\gamma \subseteq \mathcal{L}$.

We will now show that the invariant $\|y_i\|_2 \leq r$ is maintained at the end of every iteration. This suffices to prove that $x_i - y_i \in \mathbf{B}_n(0, 3\gamma + 1)$ because $x_i \in \mathbf{B}_n(0, \gamma)$ and $\|y_i\|_2 \leq 2\gamma + 1$ by the loop termination condition.

Initially, $y_i = \sum_{j=1}^{n} \alpha_j b_j$ for some coefficients α_j satisfying $0 \leq \alpha_j < 1$. Thus $\|y\|_2 \leq \sum_j \|b_j\|_2 \leq n \max_j \|b_j\|_2$, the initial value of r. Consider now the result of the change $y_i \to y_i + x_{\eta_i} - y_{\eta_i}$. We have that $\|y_i + x_{\eta_i} - y_{\eta_i}\|_2 \leq \|y_i - y_{\eta_i}\|_2 + \|x_{\eta_i}\|_2$. The first of these terms is bounded by $r/2$ because of choice of η_i in Lemma A.1. From $\|x_i\|_2 \leq \gamma$, we get that $\|y_i\|_2 \leq r/2 + \gamma$. Since the value of r gets updated appropriately, the invariant $\|y_i\|_2 \leq r$ is maintained at the end of the loop. □

The following crucial lemma says that Y_γ can be used to compute all *interesting* vectors:

Lemma A.4. *Let $v \in \mathcal{L}$ be a lattice vector such that $(2/3) \cdot \|v\|_2 \leq \gamma < \|v\|_2$. Then, with probability at least $1 - 1/2^{O(n)}$, $\exists w \in \mathcal{L}$ such that Y_γ contains both w and $w \pm v$.*

Using this lemma, we can prove our main theorem, which we restate from Section 2.1:

Theorem. (Theorem 2.1)
Given a lattice basis $U \in \mathbb{Z}^{n \times d}$, Algorithm 2.1 returns the shortest l_∞ vector in the lattice of U, with probability at least $1 - 1/2^{O(n)}$, using $2^{O(n \log n)} \cdot \|U\|^{O(1)}$ bit operations.

Proof. Define B_k to be the basis B set on Step 4 at iteration k through the for loop on line 3. For correctness, consider the iteration k such that the lattice $\mathcal{L}$ of B_k satisfies $2 \leq s(\mathcal{L}) < 3$, which we know must exist from the discussion above.

Denote by v_∞ the shortest nonzero vector in $\mathcal{L}$ under the l_∞ norm. We have that $l_2(v_\infty) \leq \sqrt{n} \cdot l_\infty(v_\infty) \leq \sqrt{n} \cdot l_\infty(v) \leq \sqrt{n} \cdot l_2(v)$, for any nonzero vector $v \in \mathcal{L}$. Hence, the l_2 norm of the shortest l_∞ vector is at most $\sqrt{n}$ times the l_2 norm of the shortest l_2 vector.

Since the length $s(\mathcal{L})$ of the shortest l_2 vector is assumed to satisfy $2 \leq s(\mathcal{L}) < 3$, we have that the l_2 norm of v_∞ satisfies $\|v_\infty\|_2 < 3\sqrt{n}$. Therefore at least one iteration of the while loop on line 7 has $(2/3) \cdot \|v_\infty\|_2 \leq \gamma < \|v_\infty\|_2$, and by Lemma A.4, with high probability some Y_γ contains w and $w \pm v_\infty$ for some $w \in \mathcal{L}$. Since the algorithm computes the differences of the vectors in Y_γ, it sets v_k to v_∞ on Step 24 with high probability.

For the cost analysis, consider a single iteration of the while loop on line 7. The value of r_0 is bounded by $(n \cdot \|U\|)^{O(1)}$. The value of m is bounded by $2^{O(n \log \gamma)} \log r_0$, which is in turn bounded by $2^{O(n \log n)} \cdot \|U\|^{O(1)}$ because $\gamma \in O(\sqrt{n})$. Since the number of sieving steps is $O(\log r_0) \in O(m)$, the total cost of a single iteration of the while loop is $m^{O(1)}$. The total number of iterations of the while loop is $O(\log n) \in O(m)$, and there are exactly $2n \in O(m)$ iterations of the outer for loop. Each arithmetic operation costs $(n \cdot \|U\|)^{O(1)} \in O(m)$, so the total cost is $m^{O(1)}$, which gives the stated bound. □

To prove Lemma A.4, a probabilistic argument will be employed. The proof can be broken into three steps. First, we identify a set of *good* points from the sampled points, and argue that this set is large. Next, we argue that there must exist a lattice point which corresponds to numerous *good* points. Finally, we argue that an imaginary probabilistic step does not essentially change the behaviour of the algorithm. Combined with the existence of a lattice point corresponding to many *good* points, this imaginary step allows us to argue that the algorithm encounters both w and $w \pm v$ for an appropriate *interesting* v.

Let v be an *interesting* lattice vector. That is, $(2/3) \cdot d \leq \gamma < d$ for $d = \|v\|_2$. For the iteration where the algorithm uses a value of γ in this range, we will denote by C_1 the points in the set $\mathbf{B}_n(v, \gamma) \cap \mathbf{B}_n(0, \gamma)$. Similarly, $C_2 = \mathbf{B}_n(-v, \gamma) \cap \mathbf{B}_n(0, \gamma)$. By choice of γ, C_1 and C_2 are disjoint. We will call the points in $C_1 \cup C_2$ *good*. The following lemma shows that probability of sampling a *good* point is large.

Lemma A.5. $\mathbf{Pr}[x_i \in C_1] \geq 2^{-2n}$.

Proof. The radius of both $\mathbf{B}_n(0, \gamma)$ and $\mathbf{B}_n(v, \gamma)$ is γ. The distance between the centers is $d = \|v\|_2$. Thus the intersection contains a sphere of radius $\gamma - d/2$ whose volume gives a lower bound on the volume of C_1. Comparing with the volume of $\mathbf{B}_n(0, \gamma)$ and using the fact that $\gamma \geq (2/3) \cdot d$, we get that

$$\mathbf{Pr}[x_i \in C_1] \geq \frac{\mathrm{Vol}(\mathbf{B}_n(0, \gamma - d/2))}{\mathrm{Vol}(\mathbf{B}_n(0, \gamma))} \geq \left(\frac{\gamma/4}{\gamma}\right)^n = 2^{-2n}.$$

□

Informally, the following lemma says that if S is large at the end of the inner loop, the set $\{x_i - y_i\}$ has many repetitions and hence is never very large.

Lemma A.6. $|Y_\gamma| \le (3\gamma + 2)^n$.

Proof. The points in $\mathcal{L}$ are separated by a distance of at least 2 since we assumed $s(\mathcal{L}) \ge 2$. Hence balls of radius 1 around each lattice point are pairwise disjoint. If we consider only the balls corresponding to points in Y_γ, all of them are contained in a ball of radius $3\gamma+2$ since $Y_\gamma \subseteq \mathbf{B}_n(0, 3\gamma+1)$ by Lemma A.2. Thus the total number of points in Y_γ is at most $\mathrm{Vol}(\mathbf{B}_n(0, 3\gamma + 2))/\mathrm{Vol}(\mathbf{B}_n(0, 1)) = (3\gamma + 2)^n$. □

The following lemma argues that there must be a lattice point corresponding to many *good* points.

Lemma A.7. *With high probability, there exists $w \in Y_\gamma$ and $I \subseteq S$ such that $|I| \ge 2^{3n}$, and for all $i \in I$, $x_i \in C_1 \cup C_2$ and $w = x_i - y_i$.*

Proof. Since $\mathbf{Pr}[x_i \in C_1 \cup C_2]$ is at least 2^{-2n} by Lemma A.5, and the number of points sampled is $\left\lceil 2^{(7+\lceil \log(\gamma) \rceil)n} \log r_0 \right\rceil$, the expected number of *good* points sampled at the start is at least $2^{(5+\lceil \log(\gamma) \rceil)n} \log r_0$. The loop performs $\log r_0$ iterations removing (by Lemma A.1) at most 5^n points per iteration. The total number of *good* points remaining in S after the sieving steps is $(2^{(5+\lceil \log(\gamma) \rceil)n} - 5^n) \log r_0 \ge 2^{(2+\lceil \log(\gamma) \rceil)n} \log r_0$ since $5^n \le 2^{3n}$.

By Lemma A.6, $|Y_\gamma| \le (3\gamma + 2)^n$. Since $3\gamma + 2 \le 4\gamma$ for $\gamma \ge (3/2)^2$, $|Y_\gamma| \le 2^{(2+\log(\gamma))n}$. Hence, there exists a $w \in Y_\gamma$ corresponding to at least $2^{(4+\lceil \log(\gamma) \rceil)n} \log r_0 / 2^{(2+\log(\gamma))n} \ge 2^{3n}$ *good* points. □

The final step in the analysis is to argue that for such a $w \in Y_\gamma$, we must also have that $w \pm v \in Y_\gamma$ with high probability for an *interesting* $v \in \mathcal{L}$.

Proof of Lemma A.4

Consider the iteration where γ satisfies $(2/3) \cdot \|v\| \le \gamma < \|v\|$ for an *interesting* lattice vector v.

It can be easily seen that $x \in C_1$ if and only if $x - v \in C_2$. Consider an imaginary process performed just after sampling all the x_i. For each $x_i \in C_1$, with probability $1/2$, we replace it with $x - v \in C_2$. Similarly, for each $x \in C_2$, we replace it with $x + v \in C_1$. (This process cannot be performed realistically without knowing v, and is just an analysis tool.) The definition of y_i is invariant under addition of lattice vectors $v \in \mathcal{L}$ to x_i, and hence the y_i remain the same after this process.

Since the sampling was done from the uniform distribution and since $(x \in C_1) \leftrightarrow (x - v \in C_2)$ is a bijection, this process does not change the sampling distribution.

We may postpone the probabilistic transformation $x_i \leftrightarrow (x_i - v)$ to the time when it actually makes a difference. That is, just before the first time when x_i is used by the algorithm. The algorithm uses x_i in two places. For $i \in J$ during

the *sieving* step, we perfom this transformation immediately after computation of J. Another place where x_i is used is the computation of Y_γ. We perform this transformation just before this computation.

In the original algorithm (without the imaginary process), by Lemma A.7, there exists a point $w \in Y_\gamma$ corresponding to at least 2^{3n} *good* points. Let $\{x_i\}$ be this large set of *good* points. With high probability, there will be many x_i which remain unchanged, and also many x_i which get transformed into $x_i \pm v$. Thus, Y_γ contains both w and $w \pm v$ with high probability.